

METHOD AND SYSTEM FOR CREATING PROFILING INDICES

Abstract of the Disclosure

A database engine and optimizer framework support creation of a series of profiling indices over a column having character string data, such as a traditional “varchar” data type. The profiling indices result in a reduction of the number of records that are searched when searching for a sub-string match within that column. In some embodiments, the series of indices are created over a column that is typically searched using the LIKE predicate or some similar technique; these indices indicate for each record whether certain sub-strings may exist in that record’s value in the column. Thus, the indices are used to find the rows that may match one or more portions of the particular term being queried or, in other words, eliminate those records that do not have at least a portion of the term to be matched. The number of records actually retrieved and searched for the query sub-string is thereby reduced.